



Cadence's Brad Griffin Digs Deep Into DDR

Guest Editor Kelly Dack stopped by the Cadence Design Systems booth at DesignCon 2015, where he sat down with Product Marketing Manager Brad Griffin to discuss Cadence's advanced PCB design and signal integrity tools, and the company's focus on DDR.

Kelly Dack: Brad, since you're the product marketing director for Cadence Design Systems, I'd like to ask a few questions about your DDR products. But first, please give us a brief overview of DDR.

Brad Griffin: I'd be happy to. One of the main points about a computer is that it has memory and you can store data in that memory—that's kind of what makes it a computing device. So they've been finding ways over the life of electronics to store and retrieve data faster out of memory. Somewhere around 2002, we came up with this idea of doubling the data rate in DDR memory, or double data rate memory.

That was unique because basically, we clocked the data into the memory, both on the rising edge and on the falling edge of the clock. It was a clever way with the same sort of signaling to basically double the data rate speeds.

KD: Was there an organization involved? Was it standardized?

BG: That's really good question. As of right now, there's a standard committee called JEDEC, and I'm going to assume they were in place back in the 2002 timeframe, but I'd have to go back and check. But obviously there's memory companies and they have to be able to plug-and-play with different controllers as they're driving the memory, so there's probably always been a standard they've been marching toward. That process used to be a lot simpler. You'd be transferring data at maybe 100 megabits per second. You would send the data, clock it in, and it wasn't nearly as complicated as it is now.

CADENCE'S BRAD GRIFFIN DIGS DEEP INTO DDR *continues*

KD: So where has DDR come from, and where is it now?

BG: There was DDR2 and then DDR3, and probably 2015 is going to be the transition where most DDR3 designs go over to DDR4. Typically, this happens because the DDR4 memory will actually become less expensive than some of the DDR3 memory.

KD: What does that mean as far as the technology from a power standpoint as well as a data standpoint?

BG: The main difference from a technology standpoint from DDR3 to DDR4 is the speed. It basically just gets faster. So any application you have in the computer that's run with DDR4 memory will make for a faster computer than one running with DDR3. One of the exciting things that has migrated probably over the last five to seven years is this new version of DDR called LPDDR, which stands for low power. That's been something primarily used in mobile devices because you certainly don't want your cell phone to run out of power in the middle of the day.

KD: With this reference to power, if I understand correctly, DDR came from a 2.5 V system and shrunk to 1.8 V and 1.5 V, and DDR4 is down at a little over 1 V. That seems really low already, so where will the LPDDR take us?

BG: If you can believe it, the LPDDR4 specification only has a 300 mV swing, so it's really low. That means that for signal integrity and power integrity engineers, there's really very little margin left. We said there was very little margin left when it was 1.5 V, and now we're down to 300 mV; this very small swing of data means that your signals have to clean and your power planes have to basically be stable. Because then you have to have a power/ground bounce associated with simultaneous switching signals. It's going to basically make



it so that you're not going to meet the signal quality requirements that JEDEC puts in place for LPDDR4. So designs are getting really interesting and what we're excited about this year at DesignCon are the things we've been putting into our tools to enable designers to be able to validate that they've done everything they need to do to meet those LPDDR4 requirements.

KD: Let's talk about your tools. Would you give us an overview of some of the advanced tools at Cadence and how you're helping designers to solve some of these higher-speed, lower-power issues?

BG: Thank you for giving me the opportunity to talk about that because we're really excited about our products. Really, the foundation for the PCB and IC package design technology at Cadence is Allegro. Allegro has been around a long time; it was called Valid a long time ago before Cadence acquired it. So that's been the place where all the actual physical implementation takes place. What we did is layer signal integrity and power integrity analysis tools on top of Allegro, which have been in place since the mid-1990s. They've been serving the market fairly well but a very exciting thing happened in 2012. Cadence acquired a company called Sigrity. Sigrity is well-known for power integrity technology and their tools called PowerDC and PowerSI, which enable both AC and DC power integrity analysis

When you merge that together with their signal integrity analysis technology, what we've been able to do is take state-of-the-art, world-class signal integrity and power integrity technology in 2012 and spent the last two and half years not only improving that technology but tightly integrating it with Allegro. Now the Allegro user base has grown accustomed to having tools where they can have signal integrity analysis on-the-fly right from the board. We're giving them advanced technology that allows them to run more advanced field solvers, more advanced analysis engines and it might not sound like that much but when we go back to the idea that we only have that 300 mV swing in LPDDR4 an integrated solution is key to converging on a working solution.

We've got this advanced analysis technology tightly integrated with the implementation environment because what will typically happen is you'll run an analysis and it doesn't work—it failed the JEDEC requirements. So, what do I have to do? I have to start working with my power plane, working with the signaling, cleaning up everything, maybe there are too many vias on the signal, etc. But once you do all this you rerun the analysis and see that you're getting closer and you start to see yourself improving. Because it's so tightly integrated, our customers can accelerate the process of finding the problem, fixing the problem and verifying you fixed the problem. It's been an exciting ride the last two and half years with Sigrity and Cadence, and the Sigrity 2015 release coming out during DesignCon is really the culmination of bringing the latest and greatest technology to the market and has addressed these very difficult design and analysis challenges around LPDDR4.

KD: We have engineers and layout people—people that specifically do SI work. Are these tools used in a team application?

BG: The challenge has been that historically, data has just been thrown over the wall: I'm the designer and I throw it over the wall; the SI guy says "fix this" and throws it back over the wall, and it's a typical back-and-forth. It's very difficult to converge. On the other hand, the work that the signal integrity and power integrity engineer performs comes from a level of expertise in his area that you can't really expect a PCB designer to have. On the other hand, the person doing integrity analysis doesn't really have the level of expertise to make the changes to the physical design that the PCB designer has. We recognize that, yet we try to provide an environment which allows the gap to be bridged as much as possible. So our Allegro PCB analysis tools, as I mentioned, have the signal integrity tools residing right on top, so we have an environment where the layout person with some level of knowledge—maybe he knows how to get IBIS model on the web and can attach that to one of his components and can make sure that all of his resistors and capacitors have proper values associated with the design database—

can actually say let me analyze the signal and see what it looks like. He may not have the expertise to know exactly how to fix it, but at least he can identify there's a problem and then just needs to determine how to resolve it.

Our approach here is that we try to let the layout person with some level of electrical background go as far as he can and then bring the expert into the same environment. It's how we've sort of structured our technology—we've got the base signal integrity technology that probably both expert and non-expert can use, and then we have advanced analysis technology that sits on top of that. The expert can go in and run the DDR simultaneous switching noise analysis. It can figure out that he's going to have 64 bits simultaneously switching and the signal is not going to work. He'll have to make some changes to the power plane to make sure it's more stable, perhaps by adding some more decoupling. He could actually with some level of expertise or knowledge of how to place things around the board put down his own capacitors, try out how it'll work and improve the overall process.

KD: Is what you're describing a radical change to front-end design with these new speeds, where it's not your classic front-end design anymore with a simple schematic passed down to a layout designer?

BG: It's an excellent question because for quite a long time Cadence has pushed what we've called a constraint-driven flow, where you do a lot of analysis upfront, create constraints, drive those constraints into design and push that forward to layout and verify it at the end. That's basically our methodology that Cadence has put in place for signal integrity, but one of the things we're showing in our booth is that we're moving this constraint-driven flow so it's not just signal integrity, but also power integrity. Because we believe that if the hardware engineer that is doing the schematic knows this



CADENCE'S BRAD GRIFFIN DIGS DEEP INTO DDR *continues*

component needs a certain amount of decoupling associated with it, instead of just putting all the decoupling capacitors on a page at the end, he basically associates his decoupling capacitors with that component in the schematic. Then when it gets to the point that you're doing actual placement of the decoupling capacitors you're going to get violations that tell you that you haven't placed the right capacitors within the right radius of this component. So we're bringing this constraint-driven flow to power integrity that's always been there for signal integrity.



KD: Excellent. Let's talk about serial interfaces. Tell us where they've come from and where we're going with them.

BG: One of the most interesting things in signal integrity is around the serial interfaces and it also sort of mixes with memory interface design as well, which is a parallel bus. With serial interfaces, the way that we typically check compliance on them is by running many signals which we call high-capacity simulation, and by many I mean like millions and tens of millions of bits. We're looking to see how many of those bits actually get transferred correctly. So when you go to the PCI-SIG, the special-interest group, they have a bit error rate test that they do with hardware. Well we can do the same sort of bit error rate testing with software. Our signal integrity software supports a high-capacity simulation and then lets you look at the eye diagram and just like with PCI-SIG we have that compliance test built into our software.

There are a couple of really interesting things about what's happening in this space. One is the most popular serial interface by far, which is PCI express. We've been at PCIe 3.0 for a few years now, and that's an 8 Gb/s interface. Most people here at DesignCon are talking about up to 28 or 56 Gb/s, so 8 is a little bit behind the bleeding edge at this point. But what's going to be happening this year is PCI-SIG is going

to approve the 4.0 spec, which is moving it to 16. Still maybe not on the bleeding edge, but doubling the data rate is very significant. One of the cool things we're showing in our booth is if someone who is using 8 Gb/s today wants to see if their same hardware will support a 16 Gb/s data transfer, we can help them check that feasibility. It's really quite interesting because you can see by default the answer is probably no, the eye is going to be closed and you're not going to meet your bit error rate testing. But because these transceivers and receivers have such advanced equalization in them we have what's called algorithmic models that sit on both sides, transmitter and receiver, and this is the same type of stuff we're going to see in devices that come out and support PCIe 4.0. We can turn on a level of equalization and see if when we boost that signal if we can open up that eye and see if it's going to meet those compliance requirements that are going to be associated with doubling the data rate from 8 to 16.

That's a pretty interesting thing that's going to be happening in 2015. But one other thing I just wanted to throw in is when we talked about LPDDR4, that data rate is actually going to go up to as high as 4266, so that's going to be working in a similar way that serial links were working about two or three years ago. The same equalization that you needed in serial links a few years ago are going to be needed in memory interfaces this year. We will support that with our algorithmic modeling interface. We can actually show today AMI modeling associated with DDR4 and LPDDR4 as well as, of course, serial links. It's just tremendously interesting and exciting with all this different technology that we get data passed across the ether into the cloud as fast as possible. All this stuff is really exciting, and the fact that we're able to analyze this and help customers get to market right the first time is what we're really excited about at Cadence, and the Allegro technology is providing that link to getting the product done right the first time.

KD: Thanks for taking the time to talk with me today, Brad.

BG: Thank you, Kelly. **PCBDESIGN**